Compressed sensing for Hamiltonian reconstruction

Kenneth Rudinger* and Robert Joynt
University of Wisconsin-Madison, Madison, Wisconsin 53706

In engineered quantum systems, the Hamiltonian is often not completely known and needs to be determined experimentally with accuracy and efficiency. We show that this may be done at temperatures that are greater than the characteristic interaction energies, but not too much greater. The condition for this is that there are not too many interactions: the Hamiltonian is sparse in a well-defined sense. The protocol that accomplishes this is related to compressed sensing methods of classical signal processing; in this case applied to sparse rather than low-rank matrices.

PACS numbers: 03.65.Wj,03.67.Lx,75.10.Dg

arXiv:1410.3029v1 [quant-ph] 11 Oct 2014

INTRODUCTION

In quantum physics, the standard method for understanding a large system has long been to make an approximate model Hamiltonian that captures the essential physics of the material in question. More recently, this situation is often turned on its head - a quantum system of n qubits is constructed and we need to find its Hamiltonian from experimental data. To do quantum information processing of any kind, accurate control of the Hamiltonian is always a prerequisite. One needs to be able to apply external controls to guide the desired time-dependent Hamiltonian, but it is usually also the case that there are "always-on" terms, generally time-independent or nearly so, in the Hamiltonian that need to be determined at a precise quantitative level [1]. This is a particularly pressing issue for a quantum memory, or in cold-atom systems that are specifically constructed in order to simulate many-body Hamiltonians. For electron spin qubits in semiconductor quantum dots [2], for example, the single-qubit energy-level splittings are subject to unknown random hyperfine fields, and there are also dipole-dipole interactions. These are one- and two-qubit interactions, but there is also the more challenging case of multi-qubit interactions. In this paper we propose an efficient way to determine these "always-on" terms.

For $n = 1$ and $n = 2$, considerable work has been done, since these cases are relevant to the performance of gates [3–5]. Process tomography is the usual tool for problems with $n > 2$, but standard methods [6, 7] require a number of measurements that scales exponentially with n. Other methods that pertain particularly to spin systems require only a small number of measurements, but they appear to involve full simulation of the system, a task that again scales exponentially [8–11]. Several authors have investigated the use of techniques from compressed sensing [12] which would give an efficient solution to this problem when the process matrix χ is $s-$sparse (has only s nonzero elements) in some basis [13–15]. The number of measurements needed to determine χ is then $O(sn)$. However, this scheme requires prior knowledge of the basis in which χ is sparse. Thus it is useful for verifying quantum gates, but cannot be used to determine entirely unknown processes (or Hamiltonians), which is the case we are considering.

As pointed out in Ref. [16], it makes sense to take advantage of the fact that, to a very good approximation, almost all qubit Hamiltonians H have only one- and two-qubit interactions, so that the number of parameters to be determined scales only as n^2. These authors suggest a sequence of randomly chosen measurements on randomly prepared states. If the time interval t between preparation and measurement is short enough: $||H||\ t << 1$, then the density matrix is simply related to H. Here $||H||$ is the operator norm (largest eigenvalue) of H. Compressed-sensing techniques can then come into play and the number of measurements required to determine H is $O(n^3)$. However, $||H||$ grows with the size of the system, which limits the usefulness of this scheme.

METHOD

Here we propose a different approach for the experimental determination of H. The most general Hamiltonian for an array of n qubits is:

$$H = -\eta \sum_{a=1}^{4^n-1} J_a \lambda_a \tag{1}$$

where a is an n-digit base-4 number $a = a_1 a_2 ... a_n$ and the λ_a are tensor products of Pauli matrices: $\lambda_a = \sigma_{a_1} \otimes \sigma_{a_2} \otimes ... \otimes \sigma_{a_n}$. $\sigma_{1,2,3} = \sigma_{x,y,z}$ and σ_0 is the identity matrix. For notational convenience we have defined the energy scale η, set by the condition that the dimensionless variables J_a satisfy $|J_a| \leq 1$. We will assume that only s of the $4^n - 1$ possible J_a are zero and $s << 4^n$. The system is placed in a bath and comes to thermal equilibrium. The density matrix is $\rho = \exp(-\beta H)/Q$, where Q is the partition function: $Q = \mathrm{Tr}\, \exp(-\beta H)$ and $\beta = 1/k_B T$. If $T = 0$, ρ reduces to $\rho = |0\rangle \langle 0|$ where $|0\rangle$ is the ground state so that the density matrix has rank 1. We will work in the opposite, high-temperature,

limit $\eta\beta << 1$, where $\rho = I - \beta H + \beta^2 H^2/2 + ...$ and we may truncate the expansion. In general there are a macroscopic number of energy eigenstates that enter ρ and ρ represents a high rank state. It is important to note that the application of compressed sensing proposed here is opposite to others in the literature that primarily focus on the determination of states of low rank [17, 18]. In fact the density matrix is technically of *full* rank at any finite temperature and the naive (but inefficient) procedure to determine the J_a would be to measure the observables λ_a. For $\eta\beta << 1$ this gives $\eta J_a = -2^{-n}\ \mathrm{Tr}\ (\lambda_a H) \approx \beta^{-1}\ \mathrm{Tr}\ (\lambda_a \rho)$. However, most of the diagonal matrix elements are exponentially small, and we will use this fact to reduce the number of measurements that need to be made.

The measurement and processing protocol is as follows. After the system reaches equilibrium, its state is given by $\rho = 2^{-n} I + 2^{-n} \sum_{a=1}^{4^n-1} v_a \lambda_a$, where $\vec{v}$ is the equilibrium polarization vector of the system. We then subject the system to a random unitary transformation U so that the new state of the system is $\rho' = U\rho U^{-1}$. The procedure for generating random U's that are efficiently implementable with a small gate set is a modification of one proposed for quantum data hiding by DiVincenzo, Leung, and Terhal [19], using work by Harrow and Low on random quantum circuits [20]. The U's are not selected uniformly from the Haar distribution but our results indicate that they provide usable compression matrices. (Details for generating each U are provided in Appendix A)

The new polarization vector $\vec{v}'$ is linearly related to the previous one: $v'_a = \sum_{b=1}^{4^n-1} C_{ab} v_b$ with $C_{ab} = 2^{-n}\ \mathrm{Tr}\left(\lambda_a U \lambda_b U^{-1}\right)$. C is an orthogonal matrix and $\vec{v}$ is a long but approximately sparse vector, the "signal vector". We now measure M of the observables λ obtaining the results $\{y_k\}_{k=1}^M$ with the y_k satisfying $-1 \le y_k \le 1$. We will discuss the magnitude of M and the choice of the λ's below. $\vec{y}$ is our "measurement vector", a subset of the elements of $\vec{v}$. We now have

$$y_k = \sum_b C^{(M)}_{kb} v_b, \tag{2}$$

where $C^{(M)}$ consists of M rows of C, the choice of rows corresponding to the measurements taken. $C^{(M)}$ is an $M \times (4^n - 1)$ matrix, the "compression matrix". The next step is to estimate the polarization vector by minimizing the L_1 norm of all possible polarization vectors that are consistent with the measurement results:

$$\vec{v}_{est} = \arg\min_{\vec{w}} ||\vec{w}||_1\ ,\ \text{subject to}\ \sum_b C^{(M)}_{kb} w_b = y_k. \tag{3}$$

The L_1 norm of a vector $\vec{w}$ is defined as $||\vec{w}||_1 = \sum_{i=1}^{d} |w_i|$. This is a convex optimization problem that can be solved efficiently. For our purposes it is important to note that this compressed sensing protocol is stable with respect to deviations from exact sparsity in the signal vector, so that, as we shall see below, the protocol works at moderate temperatures. Also, it can be shown that if $C^{(M)}$ is formed by choosing rows at random from C, then $C^{(M)}$ satisfies a certain restricted isometry condition which guarantees that that if $M > A\ n\ \ln^3 s$ we can recover $\vec{v}$ with high probability, Here A is a constant. [21].

Once a good estimate of the polarization vector is available, we can estimate the Hamiltonian:

$$H_{est} = \beta^{-1}(2^{-n}\mathrm{Tr}\left(\ln \rho_{est}\right)\ I - \ln \rho_{est}). \tag{4}$$

RESULTS

We now turn to numerical studies of the protocol for 3, 4 and 5 qubits, for which a takes on $N = 63$, $N = 255$, and $N = 1023$ values, respectively. We input a random Hamiltonian, compute the equilibrium density matrix ρ, and perform M measurements, i.e., characterize ρ by the numbers $\mathrm{Tr}(\lambda_i \rho)$, $i = 1, 2, ..., M$. (This is our definition of a measurement; while experimental measurements in a lab are subject to finite sample error, it is known that compressed sensing is robust against such errors [12].) While measurements are chosen at random, they are ordered by weight, that is all measurements of weight one (i.e., single-qubit measurements) are performed before all measurements of weight two (i.e., two-qubit measurements), and so on. (See Appendix B for further explanation.)

The simplest case is the determination of the J_a when we are given that only s of them are nonzero. We do not have firm guarantees of success at finite temperature, since the density matrix is not s-sparse. So the first task is to determine how high the temperature needs to be to ensure success. The temperature is quantified by the dimensionless ratio $\eta\beta$. Success is measured by the distance of H_{est}, the Hamiltonian estimated from Eq. 4, from the actual Hamiltonian H, the metric chosen as the one corresponding to the Frobenius norm: if $(||H_{est} - H||_F)/\eta <$ threshold, the procedure is judged to have succeeded.

Fig. 1 shows the quality of the reconstruction of H as a function of the parameters M/N, which is the number of measurements divided by the signal length, and the sparsity ratio s/N. There are 3 qubits and each pixel in the plots is the result of 100 trials. Note first that the lower right corner is a region where the number of nonzero entries in J_a is greater than the number of measurements: reconstruction is impossible there. As we move away from the diagonal to the upper left, the success probability increases. As is generally observed in cases where compressed sensing works, the boundary between success and failure (that is, the Donoho-Tanner phase transition) is sharp. High temperature is favorable for reconstruc-

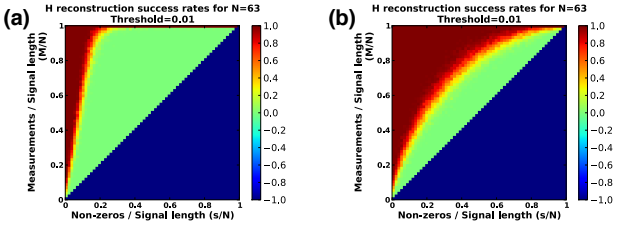

FIG. 1. Quality of Hamiltonian determination for random couplings as a function of temperature. In (a) and (b) the inverse dimensionless temperature is given by $\eta\beta = 10^{-4}$ and $\eta\beta = 10^{-1}$ respectively. Red indicates a high success rate, green indicates failure, and a "negative" success rate (blue) means that reconstruction is impossible. Each pixel is an average of 100 trials.

tion, but even at quite moderate temperatures there is a very substantial region of parameter space where the determination of H succeeds. The red region in both panels is where H is successfully reconstructed, due to the density matrix being approximately sparse in that region.

These computations show that compressed sensing can work in principle, and gives strong evidence that the number of measurements needed is proportional to n, the number of qubits, rather than N, the number of possible couplings, when the Hamiltonian is sparse. However, equipped with the knowledge that H is sparse, quantum state tomography can also be carried out with a reduced number of meaurements. We next examine the question of how much advantage is actually gained in practice over the straightforward method of standard tomography, stopping when H has been determined. Fig. 2 gives this comparison for $n = 3$ [Fig. 2(a)], $n = 4$ [Fig. 2(b)], and $n = 5$ [Fig. 2(c)], with small values of s, and for a moderate temperature of $\eta\beta = 10^{-1}$. The number of trials per data point is 100. The sampled M's have a spacing of 1 for $n = 3$ and $n = 4$, starting at a value of $M = 2$; due to computational constraints, every tenth value of M is used for $n = 5$, starting at a value of $M = 11$. The median value of the normalized quality $(||H_{est} - H||_F)/||H||_F$ of the estimate is plotted as a function of M, so that low values correspond to accurate estimates. When the curve drops off sharply, the "phase transition" from failure to success has occurred. Thus for example, in Fig. 2(a), the compressed sensing (CS) protocol for $n = 3$ and $s = 1$ succeeds at $M = 5$. It is seen that compressed sensing gives a large saving in the number of measurements for all cases considered, ranging (roughly) from a factor of 4 to 7 for $n = 3$, from 6 to 12 for $n = 4$, and from 12 to as high as 50 for $n = 5$. This is good evidence that the advantage of the compressed sensing protocol increases with n, as we would expect from the scaling arguments above.

In most cases of actual physical interest, we not only

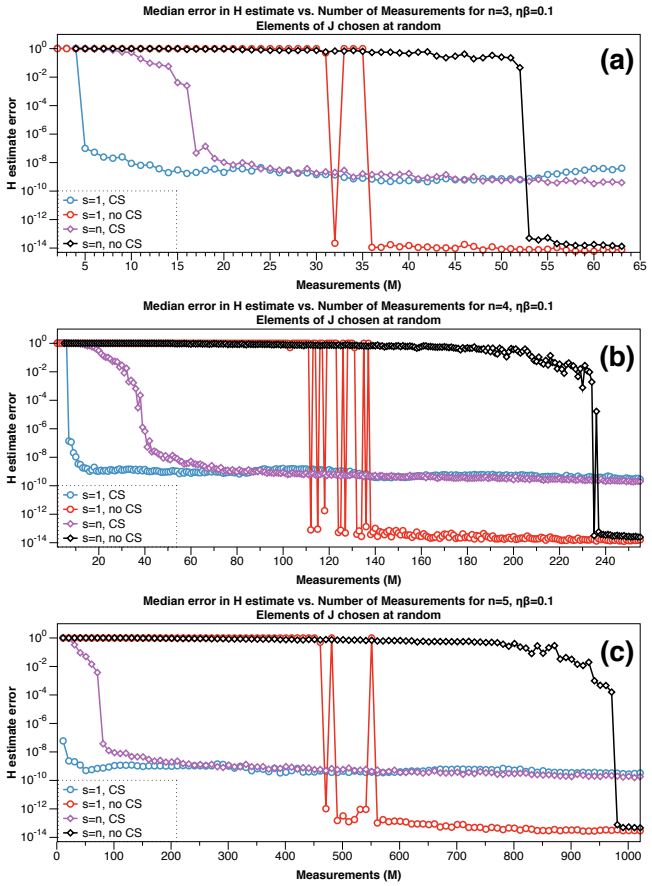

FIG. 2. Quality of Hamiltonian determination for random couplings as a function of number of qubits. (a), (b), and (c) give the error in the estimated Hamiltonian for $n = 3$, 4 and 5 qubits, respectively, as a function of the number of measurements made (M), with compressed sensing (CS) and without (no CS). [While the minimum error obtained by the no CS protocol appears to be lower than the minimum error obtained by the CS protocol ($\sim 10^{-14}$ and $\sim 10^{-9}$, respectively), these differences are simply artifacts of different noise floors of different numerical methods. For each algorithm, achieving its respective noise floor indicates that the Hamiltonian has been successfully reconstructed.] In each case, the CS protocol substantially decreases the M required to accurately reconstruct the Hamiltonian. The improvement increases with the number of qubits. Each data point is the median value of 100 trials.

have some knowledge of the sparsity of H, we also have some knowledge of where the nonzeros lie. For example, for spin qubits, 1- and 2-body interactions are likely to be much greater in magnitude than 3- and higher-body interactions. We then find $s = O\left(n^2\right)$. Locality may also reduce the sparsity; for sufficiently short-range interactions $s = O\left(n\right)$. This is a very different situation than we have considered so far, where the nonzero J_a were taken at random. Of course exponential reductions in M required to reconstruct H are now out of the question. The question is whether we can still get speedups that may be useful in real situations - even constant speedups

can be important. So we perform the same numerical experiment as in Fig. 2, but now the nonzero J_a are restricted to those corresponding to λ_a that are 1- and 2-qubit operators, *i.e.*, a has at most 2 nonzero digits. The results are shown in Fig. 3. The number of trials and all other parameters are the same as in Fig. 2. In the "no CS" (standard tomography) protocol, measurements of 1- and 2-body operators are made first, which now improves the performance of the "no CS" procedure, but not enough to overcome the advantage of the CS protocol (see Appendix C for further details).

The ratio of the number of measurements required is about a factor of 2 to 4 for $n = 3$, about a factor of 3 to 6 for $n = 4$, and about a factor of 6 to 8 for $n = 5$. Thus the speedup is less when the knowledge of the locations of the nonzeros is increased, but it is still quite substantial. More importantly, it appears that the speedup still increases with the number of qubits.

CONCLUSION

Previous improvements in efficiency of quantum state tomography have shown the usefulness of compressed sensing techniques by focusing on the reconstruction of states of low rank. This work, by contrast, uses this technique to reconstruct states of high rank. This is not useful for validation of gate quality, but it can be used to determine the parameters in a many-body Hamiltonian.

As compressed sensing reduces the number of real-valued system parameters that must be measured, at the cost of increased post-processing, compressed sensing is only of value for systems in which measurements are expensive but signal processing and post-processing are cheap. his tradeoff is highly attractive for many classical applications, but the tradeoffs vary from case to case. Our quantum protocol will be useful and attractive when measurement settings are expensive, but quantum gate operations are cheap. Otherwise, straightforward tomography will be better. The competition between the two is greatly affected by how much advance knowledge we have about the system. It is when we do not have a very good idea in advance about the shape of the Hamiltonian that our method is useful.

Acknowledgments

We thank John Aidun, Eric Bach, Charles Baldwin, Robin Blume-Kohout, S. N. Coppersmith, Daniel Crow, Adam Frees, Mark Friesen, John Gamble, Kevin Jamieson, Amir Kalev, and Robert Nowak for useful discussions.

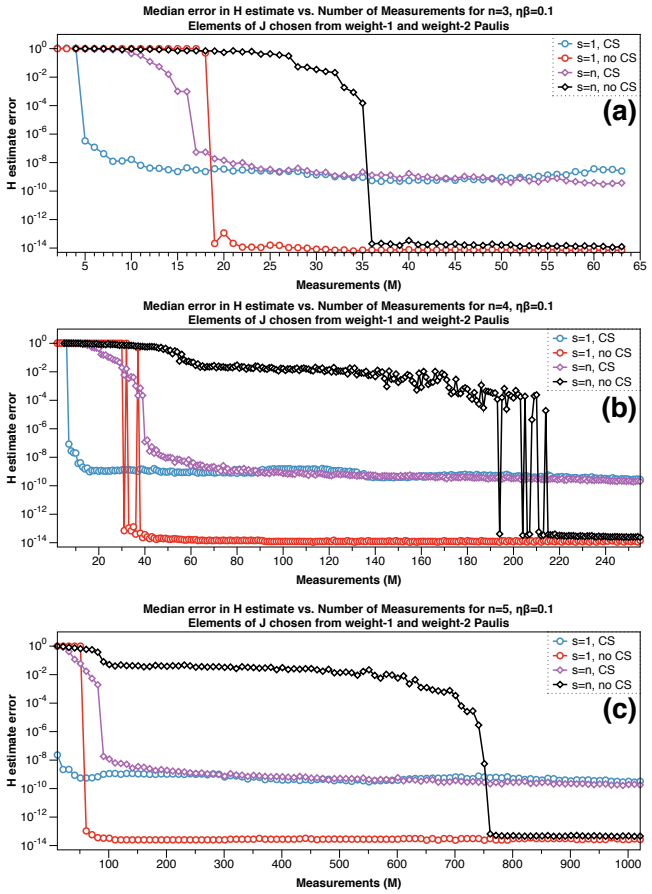
Median error in H estimate vs. Number of Measurements for n=3, ηβ=0.1
Elements of J chosen from weight-1 and weight-2 Paulis
s=1, CS
s=1, no CS
s=n, CS
s=n, no CS
(a)
H estimate error
Measurements (M)
Median error in H estimate vs. Number of Measurements for n=4, ηβ=0.1
Elements of J chosen from weight-1 and weight-2 Paulis
(b)
Median error in H estimate vs. Number of Measurements for n=5, ηβ=0.1
Elements of J chosen from weight-1 and weight-2 Paulis
(c)

FIG. 3. Quality of Hamiltonian determination for 1- and 2-qubit couplings as a function of number of qubits. (a), (b), and (c) give the error in the estimated Hamiltonian for $n = 3$, 4, and 5 qubits, respectively, as a function of the number of measurements made (M), with compressed sensing (CS) and without (no CS). (As in Figure 2, the minimum error obtained by the two different protocols is an artifact of the different noise floors for the numerical methods employed; in each cases, achieving the noise floor indicates successful Hamiltonian reconstruction.) In each case, the CS protocol substantially decreases the M required to accurately reconstruct the Hamiltonian, though this improvement is not as big as for random couplings. As before, the improvement increases as the number of qubits increases. Each data point is the median value of 100 trials.

APPENDICES

A. Generation of U

To choose a random unitary map that is efficiently implementable with a small gate set, we use the following procedure, inspired by a technique for quantum data hiding proposed by DiVincenzo, Leung, and Terhal [19], along with work by Harrow and Low on random quantum circuits [20].

For an n-qubit system, we consider the following set $\mathcal{G}$ of quantum gates

$$\mathcal{G} = \{H_p, P_q, P_r^\dagger, R_s\left(\tfrac{\pi}{8}\right), CNOT_{tu}\}, \tag{5}$$

where H is the Hadamard gate, P is the phase gate, $R\left(\frac{\pi}{8}\right)$ is the $\frac{\pi}{8}$ gate, and $CNOT$ is the controlled-not gate. The subscripts label the qubit (or qubits) that each gate is acting on, that is, $\mathcal{G}$ contains all single-qubit copies of $\{H, P, P^\dagger, R\left(\frac{\pi}{8}\right)\}$ and all two-qubit copies of $CNOT$.

To form the unitary map U, we simply select (with replacement) n^8 elements of $\mathcal{G}$ uniformly at random. Letting g_i denote the i^{th} selection from $\mathcal{G}$, we define U to be given by

$$U = \prod_{i=1}^{n^8} g_i. \tag{6}$$

Note that this gives us a random unitary operation on n qubits which, while not selected uniformly from the Haar distribution, is sufficiently random as to successfully generate a compression matrix which can be used for compressed sensing. Additionally, we note that it is an open question as to whether or not a smaller set of gates and/or a shorter gate sequence could yield equally successful results.

B. Weight-ordering of measurements

It may be of some benefit to the experimentalist for whom lower weight measurements are easier to perform to be able to prioritize low-weight measurements over high-weight measurements.

Therefore, we show here that the order the measurements are chosen in should not affect the accuracy of the Hamiltonian or the density matrix reconstructions, allowing for the measurements to be chosen according to weight. (That is to say, all single-qubit measurements may performed before any two-qubit measurement, which in turn may precede all three-qubit measurements, and so on.) This ordering by weight is justified in the following manner.

We note that if the k^{th} Pauli measured is λ_k, then the k^{th} element of our measurement vector $\vec{y}$ is given as

$$y_k = \mathrm{Tr}\left(\lambda_k U^\dagger \rho U\right), \tag{7}$$

where ρ is the initial density matrix and U is the random unitary map. However, due to the cyclic property of the trace, we may re-express Eq. (7) as

$$y_k = \mathrm{Tr}\left(\left(U\lambda_k U^\dagger\right)\rho\right). \tag{8}$$

That is, we may consider our k^{th} measurement to correspond to measuring the expectation value a Pauli subjected to a random unitary transformation with respect to the fixed and original density matrix. Therefore, as U effectively randomizes each λ_k, choosing them in order of their weights should not affect the reconstruction algorithm. (Indeed, we have performed numerical tests which demonstrate this.)

C. State reconstruction via "no CS" protocol

The "no CS" protocol for reconstructing the density matrix ρ is as follows. For an estimate of ρ in which M Pauli measurements are allowed, the M expectation values are input as the appropriate v_i's; the remaining v_i's are set to zero. While this estimation procedure could theoretically produce a non-physical ρ_{est} with negative eigenvalues, in practice this is not a concern as any state we are estimating has a polarization vector with a small L_2 norm, while a non-physical density matrix with one or more negative eigenvalues has a polarization vector with a large L_2 norm.

* Current address: Sandia National Laboratories, Albuquerque, NM 87185-1322.
Email: kmrudin@sandia.gov

[1] M. D. Shulman, S. P. Harvey, J. M. Nichol, S. D. Bartlett, A. C. Doherty, V. Umansky, and A. Yacoby, (2014), arXiv:1405.0485v1 [cond-mat].
[2] D. P. DiVincenzo and D. Loss, Superlatt Microstruct **23**, 419 (1998).
[3] J. H. Cole, S. G. Schirmer, A. D. Greentree, C. J. Wellard, D. K. L. Oi, and L. C. L. Hollenberg, Phys. Rev. A **71**, 062312 (2005).
[4] S. J. Devitt, J. H. Cole, and L. C. L. Hollenberg, Phys. Rev. A **73**, 052317 (2006).
[5] S. G. Schirmer and D. K. L. Oi, Phys. Rev. A **80**, 022333 (2009).
[6] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000).
[7] I. L. Chuang and M. A. Nielsen, Journal of Modern Optics **44**, 2455 (1997).
[8] D. Burgarth, K. Maruyama, and F. Nori, Phys. Rev. A **79**, 020305 (2009).
[9] D. Burgarth and K. Maruyama, New. J. Phys. **11**, 103019 (2009).
[10] D. Burgarth, K. Maruyama, and F. Nori, New. J. Phys. **13**, 013019 (2011).
[11] C. Di Franco, M. Paternostro, and M. S. Kim, Phys. Rev. Lett. **102**, 187203 (2009).
[12] E. Candès and M. Wakin, Signal Processing Magazine, IEEE **25**, 21 (2008).
[13] A. Shabani, R. L. Kosut, M. Mohseni, H. Rabitz, M. A. Broome, M. P. Almeida, A. Fedrizzi, and A. G. White, Phys. Rev. Lett. **106**, 100401 (2011).
[14] C. Baldwin, A. Kalev, and I. Deutsch, (2014), arXiv:1404.2877v2 [quant-ph].
[15] A. V. Rodionov, A. Veitia, R. Barends, J. Kelly, D. Sank, J. Wenner, J. M. Martinis, R. L. Kosut, and A. N. Ko-

rotkov, (2014), arXiv:1407.0761 [quant-ph].
[16] A. Shabani, M. Mohseni, S. Lloyd, R. L. Kosut, and H. Rabitz, Phys. Rev. A **84**, 012107 (2011).
[17] D. Gross, Y.-K. Liu, S. T. Flammia, S. Becker, and J. Eisert, Phys. Rev. Lett. **105**, 150401 (2010).
[18] S. T. Flammia, D. Gross, Y.-K. Liu, and J. Eisert, New Journal of Physics **14**, 095022 (2012).
[19] D. P. DiVincenzo, D. W. Leung, and B. M. Terhal, (2001), arXiv:quant-ph/0103098v1 [quant-ph].
[20] A. W. Harrow and R. A. Low, (2009), arXiv:0802.1919v3 [quant-ph].
[21] R. Vershynin, (2010), arXiv:1011.3027v7 [math.PR].